Computing optimal k-regret minimizing sets with top-k depth contours

Sean Chester, Alex Thomo, S. Venkatesh, and Sue Whitesides
Computer Science Department
University of Victoria
PO Box 1700 STN CSC
Victoria, Canada
{schester, sue}@uvic.ca, {thomo, venkat}@cs.uvic.ca

arXiv:1207.6329v1 [cs.DB] 26 Jul 2012

ABSTRACT

Regret minimizing sets are a very recent approach to representing a dataset D with a small subset S of representative tuples. The set S is chosen such that executing any top-1 query on S rather than D is minimally perceptible to any user. To discover an optimal regret minimizing set of a predetermined cardinality is conjectured to be a hard problem. In this paper, we generalize the problem to that of finding an optimal k-regret minimizing set, wherein the difference is computed over top-k queries, rather than top-1 queries.

We adapt known geometric ideas of top-k depth contours and the reverse top-k problem. We show that the depth contours themselves offer a means of comparing the optimality of regret minimizing sets with L_2 distance. We design an $\mathcal{O}(cn^2)$ plane sweep algorithm for two dimensions to compute an optimal regret minimizing set of cardinality c. For higher dimensions, we introduce a greedy algorithm that progresses towards increasingly optimal solutions by exploiting the transitivity of L_2 distance.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*geometrical problems and computations*

General Terms

Algorithms, Theory

Keywords

regret, representative databases, top-k, arrangement of lines, plane sweep

1. INTRODUCTION

For a user navigating a large dataset, the availability of a succinct representative subset of the data is crucial. For example, consider Table 1, a toy, but real, dataset consisting of the top eight scoring NBA players from the 2009 basketball season. A user viewing this data would typically be curious which of these eight players were "top of the class" that season. That is, he is curious which few tuples best represent the entire dataset, without his having to peruse it in entirety.

id	player name	points	rebs	steals	fouls
1	Kevin Durant	**2472**	623	112	171
2	LeBron James	2258	554	125	119
3	Dwyane Wade	2045	373	**142**	181
4	Dirk Nowitzki	2027	520	70	208
5	Kobe Bryant	1970	391	113	187
6	Carmelo Anthony	1943	454	88	225
7	Amare Stoudemire	1896	732	52	**281**
8	Zach Randolph	1681	**950**	80	226

Table 1: Statistics for the top eight NBA point scorers from the 2009 regular season, taken from `databasebasketball.com`. The top score in each statistic is bolded.

A well-established approach to representing a dataset is with the *skyline* operator which returns all pareto-optimal points.[1] The intention of the skyline operator is to reduce the dataset down to only those tuples that are guaranteed to best suit the preferences or interests of *somebody*. If the toy dataset in Table 1 consisted only of the attributes *points* and *rebounds*, then the skyline would consist only of the players *Kevin Durant*, *Amare Stoudemire*, and *Zach Randolph*, so these three players would represent well what are the most impressive combinations of point-scoring and rebounding statistics. The skyline is a powerful summary operator only on low dimensional datasets, however; even on this toy example, *everybody* is in the skyline if we consider all four attributes. In general, there is no guarantee that the skyline is an especially succinct representation of the dataset.

Regret

A promising new alternative is the *regret minimizing set*, introduced by Nanongkai et al. [14], which hybridizes the skyline operator with top-k queries. A top-k query takes as input a utility function f and evaluates each tuple according to f, reporting the k tuples with highest values. Figure 1 shows how highly three of the points rank for a user utility function of $f(\text{pts}, \text{rebs}) = (\text{pts} + \text{rebs})/2$, if the attributes are normalized. The distance from the orthogonal line of each point is proportional to the point's score for that user weight. This reveals that Randolph earns the highest normalized

[1]Pareto-optimal points are those for which no other point is higher ranked with respect to every attribute.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 2012 ACM X-XXXXX-XX-X/XX/XX ...$10.00.

score (0.840), compared to Kevin Durant (0.828) and then Kobe Bryant (0.604).

To evaluate whether a subset effectively represents the entire dataset well, Nanongkai et al. introduce *regret ratio* as the ratio of how far from the best score in the dataset is the best score in that subset, with respect to a given utility function. Graphically, this is proportional to how much smaller than the largest arrow is the largest arrow in the subset. For the subset $\{\text{Bryant}, \text{Durant}\}$, the regret ratio is:

$$(0.840 - 0.828)/0.840 = 0.0143,$$

since the score for Randolph is the best in the dataset at 0.840, and the score for Durant is the best in the subset at 0.828.

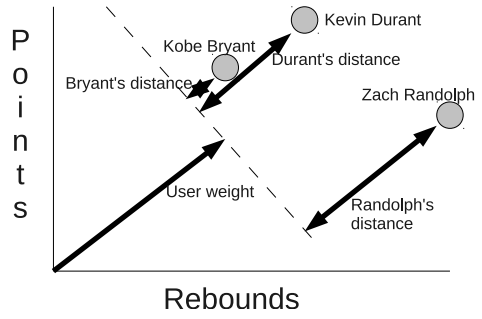


Figure 1: A utility function $f(\text{pts}, \text{rebs}) = (\text{pts} + \text{rebs})/2$ represented as a vector $\vec{f} = \langle .5, .5 \rangle$, and three data points from Table 1 shown with their scores being proportional to the distance from the line orthogonal to $\vec{f}$.

Motivated to derive a succinct representation of a dataset, one with fixed cardinality, Nanongkai et al. introduce *regret minimizing sets* [14], posing the question, "Does there exist one set of c tuples that makes every user at least x% happy?" A regret minimizing set is a subset of a dataset that minimizes the regret ratio.

A linear top-k query can be considered as a problem of projection [6], where each tuple is regarded as a vector, as is the utility function. The score of a tuple is proportional to the size of its projection onto the utility function vector and its scores for every possible utility function trace a (hyper-)sphere emanating from the point. Minimizing regret ratio is equivalent to finding a subset that minimizes the maximum distance between the "best" spheres in the subset and the "best" spheres in the entire dataset, as illustrated in Figure 2. Of the eight basketball players of Table 1, *Zach Randolph* achieves this criteria, so he is the optimal regret minimizing set of order (i.e., size) 1.

Randolph, the optimal regret minimizing tuple, however, is a peculiar choice to represent the dataset of Table 1 since he is the worst rated with respect to points. This exposes a weakness of regret minimizing sets: they are based on assuming that a "happy" user is one who obtains their absolute top choice. However, for an analyst curious to know what is a high point-scoring player, is he really dissatisfied with LeBron James as a query response rather than Kevin Durant?

To change the scenario a bit, consider a dataset of hotels and a user searching for one that suits his preferences. The absolute top theoretical choice may not suit him especially well at all. It could be fully booked. Or, he might know that the manager reminds him of his ex-wife. Regardless, it makes sense to present him a few, say k, options, any of with which he would be happy.

We generalize the concept of regret and of regret minimizing sets to that of k-regret, analogous to the difference between top-k queries and top-1 queries, because top-k is often a better threshold for "happiness". The analogous problem is to find a subset S of points in the dataset that minimize the distance from the best point in S to the k'th best point in the entire dataset. This relaxation prevents having to fit an outlier tuple like Randolph.

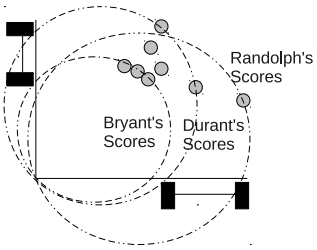


Figure 2: The spheres of scores for Durant, Bryant, and Randolph for *rebounds* as the x-attribute and *points* as the y-attribute. The regret ratio of a subset S is (roughly) the ratio of the distance from the best possible score in S from the best possible score on the entire dataset. For Durant and Bryant, this is maximized on the x-axis (where the best score is instead Randolph), and for Randolph, conversely, this is maximized on the y-axis (where the best score is instead Durant).

Optimality

A fundamental open question remains with regards to both the problem introduced by Nanongkai et al. and our generalisation of it. How can one efficiently compute the optimal k-regret minimizing set of a predetermined cardinality c, the subset that achieves the minimal regret ratio of all size c subsets of the dataset? This is a problem conjectured to be *NP*-Hard by Nanongkai et al. for $k = 1$: it involves searching for the best among $\mathcal{O}(n^c)$ different subsets.

We introduce algorithms that strive to compute optimal k-regret minimizing sets. Towards this end, we relate the recent work on top-k depth contours of Chester et al. [7] for the reverse top-k problem of Vlachou et al. [19]. The top-k depth contours are a dual space, geometric idea that succinctly represent exactly the k'th ranked tuple for all utility functions. We demonstrate that these ideas are directly applicable to discovering optimal k-regret minimizing sets. For instance, if the cardinality restriction is lifted, then the contours are *precisely* the optimal solutions (Lemma 3.1). In the presence of cardinality restrictions, the contours still aid in finding optimal solutions (Theorem 1).

1.1 Contributions and Outline

In this paper we propose the first algorithms for computing optimal k-regret minimizing sets. In particular, we:

- generalize *regret* and *regret minimizing sets*, top-1 concepts, to those of *k-regret* and *k-regret minimizing sets*, top-k concepts (Section 2);
- identify that apparently unrelated work on top-k depth contours and reverse top-k queries [7] sheds insight into the problem of identifying *optimal* k-regret minimizing sets (Section 3);
- introduce an $\mathcal{O}(n^2 c)$ algorithm to compute the optimal size-c k-regret minimizing subset S of a two-dimensional dataset D for the family of positive linear functions $\mathcal{L}^+$, despite a conjecture by Nanongkai et al. [14] that the general dimension problem is intractible for $k = 1$ (Section 4);

- introduce a greedy algorithm for general dimensions that leverages the relationship between top-k depth contours and optimality in order to progress towards more optimal solutions (Section 5); and
- relate our work within the context of other literature (Section 6).

2. PRELIMINARIES

In the following sections, we will demonstrate how to compute optimal k-regret minimizing sets by equating the problem to one in dual space. Within the dual space, the optimal solution is the top-k depth contour if it small enough, or else the convex chain through the arrangement of lines in dual space that minimizes a particular distance ratio. Before embarking on these objectives, however, we introduce some concepts formally in four subsections: first, k-regret and k-regret minimizing sets (Section 2.1); next, the transformation of the data into an arrangement of lines in dual space and some of the tools therein that we use (Section 2.2); penultimately, the top-k depth contours that exist in the arrangement of lines and are fundamentally connected to finding optimal k-regret minimizing sets (Section 2.3); and, finally, the problem definition under study (Section 2.4). Throughout the paper, we consider the family of positive linear functions $\mathcal{L}^+$, which, without loss of generality, can be reduced to the family of positive unit linear functions $\mathcal{U}^+$ [6]. Nonetheless, we present the definitions for general family.

2.1 k-Regret

k-Regret, introduced here, is a generalisation of *regret*, introduced by Nanongkai et al. [14]. We recall the definitions from that paper and introduce the generalisation in this subsection.

Given a dataset D of n d-dimensional numeric tuples, a subset $S \subseteq D$, a family of utility functions $\mathcal{F}$, and a utility function $f \in \mathcal{F}$:

DEFINITION 2.1 (GAIN [14]). *The* gain *for a subset $S \subseteq D$ on $f \in \mathcal{F}$ is:*

$$gain(S, f) = \max_{p \in S} f(p).$$

That is to say, the *gain* of a subset S, given a utility function f, is simply the highest score achievable in S for the function f. Recalling the example of Table 1 and the utility function $f(\text{pts}, \text{rebs}) = (\text{pts} + \text{rebs})/2$, and assuming the data is normalized, the gain of $\{\text{Bryant}, \text{Durant}, \text{Wade}\}$ is 0.828. The generalisation of *gain* is to *k-gain*:

DEFINITION 2.2 (k-GAIN). *Consider a descending order list of $f(p)$ for all $p \in S \subseteq D$, given $f \in \mathcal{F}$. Then, the* k-gain *of S on f is simply the k'th value in the list.*

In other words, the *k-gain* for a subset $S \subseteq D$ is the k'th best score achieved by a point in S on the utility function f. For the subset $S = \{\text{Bryant}, \text{Durant}, \text{Wade}\}$ and the same function f, the 2-gain is the second best score, 0.748, the score for *Durant*. For $k = 1$, this definition is equivalent to Definition 2.1.

Regret ratio, then, is a reflection of how well the gain of a subset approaches that of the entire dataset.

DEFINITION 2.3 (REGRET AND REGRET RATIO [14]). *The* regret *for a subset $S \subseteq D$ on $f \in \mathcal{F}$ is:*

$$r_D(S, f) = gain(D, f) - gain(S, f).$$

The regret ratio *is:*

$$rr_D(S, f) = \frac{r_D(S, f)}{gain(D, f)}.$$

Since the best score for f is 0.840 the regret for the running example S is $(0.840 - 0.828)$ and the regret ratio is $(0.840 - 0.828)/0.840$. We generalise this to k-regret by evaluating how well the gain of a subset approaches the *k-gain* of the entire dataset. Note, again, that this reduces to Definition 2.3 if $k = 1$.

DEFINITION 2.4 (k-REGRET AND k-REGRET RATIO). *The* k-regret *is:*

$$kr_D(S, f) = \max(kgain(D, f) - gain(S, f), 0).$$

The k-regret ratio *is:*

$$krr_D(S, f) = \frac{kr_D(S, f)}{kgain(D, f)}.$$

Since *Durant* is the second highest scoring tuple in D for f, the 2-regret ratio of $S = \{\text{Bryant}, \text{Durant}, \text{Wade}\}$ is $(0.828 - 0.828)/0.828 = 0$. The subset S perfectly matches the top-2 requirement for utility function f. Finally,

DEFINITION 2.5 (MAXIMUM k-REGRET RATIO). *The* maximum k-regret ratio *for a subset $S \subseteq D$ with respect to a* family *of utility functions $\mathcal{F}$ is:*

$$krr_D(S, \mathcal{F}) = \sup_{f \in \mathcal{F}} krr_D(S, f).$$

The *maximum k-regret ratio* is the largest observable k-regret ratio for any utility function in an entire family. For S the 1-regret ratio is maximized for $g(\text{pts}, \text{rebs}) = \text{rebs}$, at which the best score obtainable is S is 0.655 and the 1-regret ratio is $(1.000 - 0.655)/1.000$ and the 2-regret ratio is $(0.771 - 0.655)/0.771$.

Finally, a k-regret minimizing set of order c is simply one with cardinality c that minimizes the maximum k-regret ratio. There exist *optimal* k-regret minimizing sets of order c, which are those that achieve minimal maximum k-regret ratio of all subsets of size c.

DEFINITION 2.6 (OPTIMAL k-REGRET MINIMIZING SET). *An* optimal k-regret minimizing set *of order c on a dataset D given a family of utility functions $\mathcal{F}$ is:*

$$\mathcal{S}_c(D, \mathcal{F}) = \text{argmin}_{S \subseteq D, |S| \leq c} krr_D(S, \mathcal{F}).$$

As well, Definition 2.6 reduces to that of Nanongkai et al. [14] if $k = 1$.

2.2 Arrangements of Lines

The algorithms that we propose in this paper are geometric in nature and operate on arrangements of hyperplanes in dual space. Arrangements of hyperplanes (or lines, in two dimensions), are well studied in Computational Geometry and are induced by the intersections of a set of hyperplanes.

DEFINITION 2.7 (ARRANGEMENT). *An* arrangement *of a set of d-dimensional hyperplanes $\mathcal{H}$, denoted $\mathcal{A}_\mathcal{H}$, is a partitioning of $\mathbb{R}^d$ into cells, edges, and vertices. Each* cell *is a connected component of $\mathbb{R}^d \setminus \mathcal{H}$. Each vertex is an intersection point of some d hyperplanes in $\mathcal{H}$. An edge is a line segment between two vertices of $\mathcal{A}$.*

We arrive at an arrangement of hyperplanes by applying the duality transform introduced by Chester et al. [7], which fixes an arbitrary positive real τ and converts every point $p_i \in D$ to a hyperplane h_i by considering p as a vector $\vec{p}$ and constructing the hyperplane h_i to be all vectors $\vec{x}$ that solve $\vec{p} \cdot \vec{x} = \tau$.

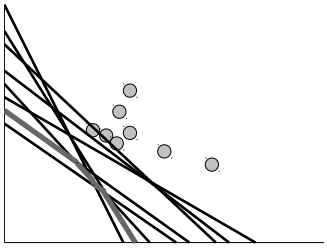


Figure 3: The eight basketball players from Table 1, considering only the attributes *points* and *rebounds*, both normalized. The tuples are transformed into translated nullspace equations, and the resulting arrangement of lines is shown. Also depicted in thicker, light magenta lines is the second top-k contour [7], a succinct representation of the 2nd ranked tuples for any top-k query.

DEFINITION 2.8 (TRANSLATED NULLSPACE TRANSFORM [6]). *Given a fixed positive real, τ, and a dataset of d-dimensional points D, the* translated nullspace transform *transform each primal space point $p_i \in D$ into a dual space (d-1)-hyperplane h_i (or line l_i in two dimensions) composed of all vector solutions to the equation $\vec{p_i} \cdot \vec{x} = \tau$.*

For the basketball example, considering only the attributes *points* and *rebounds*, which have first been normalized, the arrangement of lines produced by the translated nullspace duality transform is illustrated in Figure 3. Note that the intersection points of two lines l_i and l_j occur exactly in the direction of the vector $\vec{f}$ for which $f(p_i) = f(p_j)$.

Two other important concepts that are central ideas in Computational Geometry and of high relevance to this paper are *lower envelopes* of arrangements of lines and *convex chains* within arrangements of lines.

DEFINITION 2.9 (LOWER ENVELOPE). *The lower envelope of an arrangement of lines is the set of edges under which no other edges exist.*

For the purposes of this paper, in which we consider only the positive quadrant of Euclidean space, the *lower envelope* is the set of edges closest to the origin, O.

DEFINITION 2.10 (CONVEX CHAIN). *A convex chain in an arrangement of lines $\mathcal{A}_\mathcal{L}$ is the lower envelope in the arrangement of some subset $\mathcal{L}' \subseteq \mathcal{L}$ of lines, $\mathcal{A}_{\mathcal{L}'}$.*

Alternatively, a convex chain can be considered to be any set of edges in the arrangement that form a convex polygon with O.

2.3 Top-k Depth Contours

We also recall here two definitions to establish what are top-k depth contours, since they form the basis of our algorithms.

DEFINITION 2.11 (TOP-k DEPTH [7]). *The* top-k rank depth *of a point p within an arrangement $\mathcal{A}$, is the number of edges of $\mathcal{A}$ between p and the origin. That is to say, the depth of p is the number of intersections between edges of $\mathcal{A}$ and $[\mathcal{O}, p]$.*[2] *Similarly, the* top-k rank depth *of a cell or edge of $\mathcal{A}$ is the top-k rank depth of every point within that cell.*

[2]We remark that this is identical to the more familiar concept of a *level* if all the lines pass through the positive quadrant, as we assume here. Nonetheless, we adopt this definition because there is no reason why the techniques described in this paper cannot be extended easily to handle attributes that range into negative values.

In their paper, Chester et al. [7] show that the rank of a point in a dataset D is precisely its top-k rank depth, and that top-k rank depth creates a series of n contours in $\mathbb{R}^d$, the i'th of which is comprised of the transformed points that had rank exactly i in D.

DEFINITION 2.12 (TOP-k RANK DEPTH CONTOUR [7]). *A* top-k rank depth contour *is the set of edges in an arrangement $\mathcal{A}_\mathcal{L}$ that have top-k rank depth exactly k.*

2.4 Problem Definition

Now, we can formally describe the problem under study in this paper:

Problem Definition 1. Given any integer c and set D of n d-dimensional points, find an optimal k-regret minimizing set of order c, $\mathcal{S}_c(D,\mathcal{U}^+)$, for the family of positive unit linear functions $\mathcal{U}^+$.

3. A CONTOUR VIEW OF REGRET

In this section, we show that the concept of regret that was introduced by Nanongkai et al. [14]–and also the generalisation we introduce in this paper–are strongly connected to the dual space concept of top-k depth contours introduced by Chester et al. [7]. More precisely, we prove Theorem 1 which equates the problem of finding an optimal k-regret minimizing set to a dual space problem of finding a set of lines that are "closest" to the top-k depth contour. This alternative formulation of the problem facilitates designing algorithms in the dual space for two dimensions (Section 4) and general dimension (Section 5) to find optimal regret minimizing sets.

The argument proceeds by showing that the contour itself, $\mathcal{C}_k$, is the optimal solution, provided that it is small enough, $|\mathcal{C}_k| \leq c$ (Lemma 3.1). In the dual space, the regret ratio of a line relative to another line, given a utility function f, is given by the relative distances of the lines from the origin in the direction indicated by f (Lemma 3.2). So, the evaluation of regret in dual space is a scaled Euclidean distance computation.

We also show that the best options available to users within a set of points $S \subseteq D$ is exactly given in the the set of dual lines of S, their lower envelope (Lemma 3.4). So, minimizing the scaled distance of that envelope from the contour yields an optimal solution (Theorem 1).

LEMMA 3.1. *The set of points contributing to $\mathcal{C}_k$ is a k-regret minimizing set $\mathcal{S}_c(D, \mathcal{L}^+)$ if $|\mathcal{C}_k| \leq c$.*

PROOF. $\mathcal{C}_k$ is constructed such that, for any linear function $f \in \mathcal{L}^+$, a point p on $\mathcal{C}_k$ has rank exactly k. Therefore, $krr_D(\{p\}, f) = 0, \forall f \in \mathcal{L}^+$. □

To summarize Lemma 3.1, the contour is necessarily an optimal solution for any $c \geq |\mathcal{C}_k|$ because it is a representation of the k'th ranked tuple for any linear utility function, and the k-regret ratio of the k'th ranked tuple is 0.

Since the k-contour represents the barrier of no k-regret, any points transformed to lines farther from the origin O than the contour have positive regret proportional to the distance from O. Conversely, any points transformed to lines closer to the origin than $\mathcal{C}_k$ with respect to f have k-regret=0, because they are within the top-k on f.

LEMMA 3.2. *For any utility function $f \in \mathcal{U}^+$ and tuple $p_i \in D$ transformed to line $l_i \in \mathcal{L}$, let $\Delta_{\mathcal{C}_k}$ denote the distance of $\mathcal{C}_k$ from O with respect to $\vec{f}$, Δ_i, the distance of l_i from O, and $\Delta' \geq 0$ denote the distance of l_i from $\mathcal{C}_k$. Then, $krr_D(\{p_i\}) = \Delta'/\Delta_i$.*

PROOF. Recall that each line $l_i \in \mathcal{L}$ is constructed of vectors $\vec{x}$ such that $\vec{p_i} \cdot \vec{x} = \tau$. So, since $||\vec{f}|| = 1$, the distance of l_i from O in the direction of $\vec{f}$ is $\tau / f(p_i)$.

Thus, if $p_{\mathcal{C}_k} \in D$ represents a point on $\mathcal{C}_k$ in the direction of $\vec{f}$,

$$\begin{aligned} \Delta' &= \Delta_i - \Delta_{\mathcal{C}_k} \\ &= \frac{\tau}{f(p_i)} - \frac{\tau}{f(p_{\mathcal{C}_k})} \\ &= \tau \frac{f(p_{\mathcal{C}_k}) - f(p_i)}{f(p_i) f(p_{\mathcal{C}_k})} \\ &= \Delta_i (krr_D(\{p_i\})) \end{aligned}$$

□

COROLLARY 3.3. *For a fixed $f \in \mathcal{U}^+$, let Δ_i denote the non-negative distance for some line $l_i \in \mathcal{L}$ from $\mathcal{C}_k$ in the direction of $\vec{f}$. Then, $\Delta_i \propto krr_D(\{p_i\}, f)$.*

Lemma 3.2 establishes that the the regret for a singleton set $\{p_i\}$ on a function f in dual space is related to the Euclidean distance of the transformed line l_i to O and to $\mathcal{C}_k$. Corollary 3.3 notes that, if we consider only a single utility function, then the regret for different singleton sets can be straightforwardly compared by the distance from $\mathcal{C}_k$, since the distance of $\mathcal{C}_k$ to O is static. Next, we show that, given a non-singleton set, the maximum k-regret can be evaluated efficiently, because it is to be observed on the lower envelope in the dual space.

LEMMA 3.4. *For a set $S \subseteq D$, let $\mathcal{L}$ denote the set of lines produced by transforming every point $p_i \in S$ into its translated nullspace, l_i. The lower envelope of $\mathcal{L}$ captures the maximum gain (and, ergo, minimum regret) of S for any $f \in \mathcal{U}^+$.*

PROOF. For any $f \in \mathcal{U}^+$, the nearest line to O is that line l which is on the lower envelope in the direction of $\vec{f}$. Since l has the smallest distance to O of all lines in $\mathcal{L}$, it also has the smallest distance (possibly negative) to $\mathcal{C}_k$ of all lines in $\mathcal{L}$. By Corollary 3.3, p_i then also has the minimum regret with respect to f of all $p \in S$. □

Furthermore, while the lower envelope captures all the interestingness of a set S, there are only particular points on the envelope where the maximum regret can occur: the points where the distance ratio could be maximized. Lemma 3.5 confirms that these points are exactly the vertices of the lower envelope and the vertices of the contour.

LEMMA 3.5. *Consider a contour $\mathcal{C}_k$ and a convex chain of line segments $\mathcal{C}$. Let Δ_i denote the distance of $\mathcal{C}$ from O in the direction of $\vec{f}$ and, similarly, Δ', the distance of $\mathcal{C}$ to $\mathcal{C}_k$. The expression Δ'/Δ_i is maximized either at a vertex of $\mathcal{C}_k$ or a vertex of $\mathcal{C}$.*

PROOF. Both $\mathcal{C}_k$ and $\mathcal{C}$ are piecewise linear, so the expression Δ'/Δ_i can only be maximized at some junction point. □

Since Lemmata 3.4 and 3.5 permit evaluating regret in the dual space of translated nullspaces, we can derive an alternative view of the problem of finding an optimal k-regret minimizing set, as shown in Theorem 1 and illustrated in Figure 4.

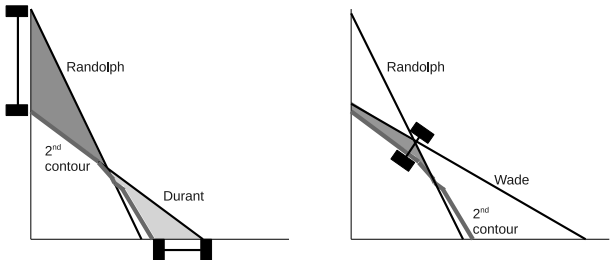


Figure 4: An illustration of 2-regret. Shown left are two different order-1 sets, {Durant} and {Randolph}, along with the second contour. The axes are the directions of maximum regret for the respective sets. To the right is shown the (non-optimal) order-2 set {Randolph, Wade}.

THEOREM 1. *Let g denote the function which transforms any point $p \in D$ to its translated nullspace and let $\mathcal{C}_k$ denote the top-k depth contour of D. The optimal k-regret minimizing set S of D with size at most c on the family of positive unit linear functions $\mathcal{U}^+$ is exactly the set of lines $\mathcal{L} = \{g(p), p \in S\}, |\mathcal{L}| \leq c$, the lower envelope $\mathcal{E}$ of which minimizes the maximum ratio of distances from $\mathcal{E}$ to $\mathcal{C}_k$ and $\mathcal{E}$ to O at any vertex of $\mathcal{E}$ and of $\mathcal{C}_k$.*

PROOF. From Lemma 3.2, the regret ratio for each point in S is equivalent to the ratio of $\mathcal{E}$ to $\mathcal{C}_k$ and $\mathcal{E}$ to O, and from Lemma 3.4, the best such ratio is on the lower envelope of $\mathcal{L}$. From Lemma 3.5, this must occur at either a vertex of $\mathcal{C}_k$ or a vertex of $\mathcal{E}$. □

A final remark is with regard to the two dimensional case, in particular. We note that any lower envelope is, in fact, a convex chain, so the two dimensional problem can be viewed rather as searching for the best convex chain.

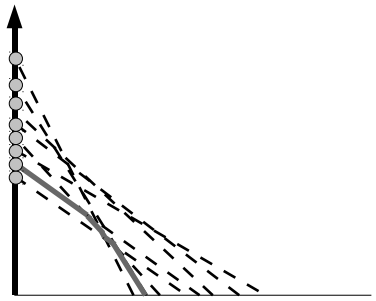
LEMMA 3.6. *Let g denote the function which transforms any point $p \in D$ to its translated nullspace and let $\mathcal{C}_k$ denote the top-k depth contour of D. The optimal k-regret minimizing set $S \subseteq D \subseteq \mathbb{R}^2$ with size at most c is exactly the convex chain $\mathcal{C}$ through the arrangment of lines $\mathcal{L} = \{g(p), p \in S\}$ that has at most $c-1$ turns and that minimizes the maximum ratio of the distance from $\mathcal{C}$ to $\mathcal{C}_k$ and the distance of $\mathcal{C}$ to O.*

PROOF. Note that any lower envelope of a set of lines is, in fact, a convex chain and that, by convexity, any line can appear at most once on the lower envelope. So, a convex chain with c lines will have $c-1$ turns. Also, any convex chain with $c-1$ turns that is optimal must be a lower envelope of some set of lines, otherwise the sequence of turns that follows the lower envelope of the same set of lines will be more optimal. By Theorem 1, this is the optimal k-regret minimizing set of size c for D. □

4. AN ALGORITHM FOR TWO DIMENSIONS

As we showed in Lemma 3.6, solving the regret minimizing problem reduces to finding the best convex chain $\mathcal{C}$ with fewer than c turns through an arrangement. The optimal solution is the one which minimizes the distance ratio of $\mathcal{C}$ to $\mathcal{C}_k$ and $\mathcal{C}$ to O. There are potentially $\binom{n}{c}$ different convex chains with at most $c-1$ turns; however, so we need to improve upon the $\mathcal{O}(n^c)$ naive algorithm which tries every combination.

We offer here a plane sweep, dynamic programming algorithm which runs $\mathcal{O}(cn^2)$ and independently of k. The algorithm follows each of the n translated nullspace lines in $\mathcal{L}$ radially from the



(a) The arrangement of lines after initialisation

$\mathcal{Q}$

Lines	*slope of r*
(Wade, Nowitzki)	35.71
(Anthony, Stoudemire)	15.39
(Bryant, Anthony)	6.07
(Stoudemire, Randolph)	2.64

(b) The priority queue, including all initial intersections

$\mathcal{P}$

Line	$m = 1$	$m = 2$
Durant	0.000	0.000
James	0.000	0.000
Wade	0.057	0.057
Nowitzki	0.062	0.062
Bryant	0.080	0.080
Anthony	0.089	0.089
Stoudemire	0.105	0.105
Randolph	0.188	0.188

(c) The initial cost values for each line, the difference in y-intercept to the contour

Figure 5: The initialisation of the algorithm for two dimensions. The sweep ray begins at the y-axis. The priority queue $\mathcal{Q}$ is populated with intersection points of lines that neighbour on the y-axis and intersect in the positive quadrant, sorted by the order in which r will pass through them. Every entry in the path list $\mathcal{P}$ is originally set to the distance of the corresponding line to the contour along the y-axis.

y-axis to the x-axis along a sweep line r, processing intersection points, and remembering c of the best paths yet encountered for each line. We maintain three data structures, each which maintains some invariant with respect to the current position of r. The optimal solution can then be read from one of the data structures at the conclusion of the radial plane sweep.

As one scans the plane, the data structures need to be updated whereever two lines intersect in order to reflect the changed state of the arrangement of lines and to maintain the invariants. All three of our data structures need to be updated at and only at line intersections.

4.1 Data Structures

As mentioned, there are three data structures. The first is the set of lines, $\mathcal{L}$, sorted by their distance from the origin in the direction of r. The second is a priority queue, $\mathcal{Q}$, containing intersection (or *event*) points yet to be processed, sorted by the order in which r will pass through them. The third data structure, $\mathcal{P}$ is for the dynamic programming component and maintains for each line the best solutions seen between the y-axis and the current position of r. The structure $\mathcal{P}$ is an n by c matrix. In each cell (i, j) is stored the optimum convex chain from r back to the y-axis which both contains at most j lines and ends on line l_i.

4.1.1 Data structure transitions

We first describe the algorithm by how the data structures evolve at each intersection point. Both $\mathcal{Q}$ and $\mathcal{L}$ behave as in traditional plane sweep algorithms, whereas $\mathcal{P}$ behaves as in a traditional dynamic programming algorithm.

$\mathcal{L}$

Consider an intersection point $p_{i,j}$, the intersection of lines l_i and l_j. Because the lines are intersecting, we know they are immediately adjacent in $\mathcal{L}$. We swap l_i and l_j in $\mathcal{L}$ to reflect the fact that immediately after $p_{i,j}$, they will have opposite order as compared to beforehand.

$\mathcal{Q}$

The priority queue contains all those intersection points that are between r and the positive x-axis that feature two lines which have been adjacent at some point between the positive y-axis and r. Again, consider an intersection point $p_{i,j}$. Immediately thereafter, lines l_i and l_j have been swapped. As such, there are potentially two new intersection points to add to $\mathcal{Q}$, namely l_i and his new neighbour (should one exist) and l_j and his new neighbour (again, should one exist). Both these intersection points are added to the appropriate place in $\mathcal{Q}$, provided that they are between r and the positive x-axis. The point $p_{i,j}$ is removed.

$\mathcal{P}$

Again, consider an intersection point $p_{i,j}$ featuring lines l_i and l_j. Let l_i be farther from the origin than l_j in the direction of a ray *after* r. There are three valid paths through $p_{i,j}$, as illustrated in Figure 6.

First consider the line, l_i, that emerges above after $p_{i,j}$ (line l_2 in Figure 6). Consider also the row of c cells of $\mathcal{P}$ that describe best paths for l_i. The h'th such cell, describes the chain with optimum distance to $\mathcal{C}_k$ that uses at most $h - 1$ turns and emerges from $p_{i,j}$ along line l_i. Because the turn (l_j, l_i) is invalid, paths for l_i cannot change, only their costs. The cost is updated to the larger of what the value was before and the distance from $p_{i,j}$ to $\mathcal{C}_k$ in the direction of r.

So, considering first a chain that leaves along l_j, the best possible route to the next intersection point of l_j is exactly whatever was

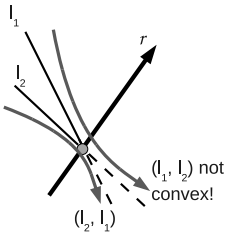
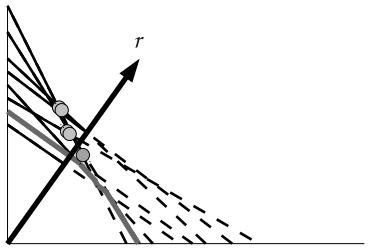


Figure 6: An illustration of the three possible paths through an intersection point. Either line l_1 or l_2 could simply pass through. Because an envelope of lines must form a convex chain, on the other hand, only l_2 has the luxury of turning onto l_1. The path (l_2, l_1) requires an "illegal" concave turn.



(a) The arrangement of lines immediately before processing the darker point, (Stoudemire, Randolph).

$\mathcal{Q}$

Lines	*slope of r*
(Stoudemire, Randolph)	2.64
(Nowitzki, Randolph)	**2.48**
(Nowitzki, Stoudemire)	2.22
(Wade, Anthony)	2.07
(James, Nowitzki)	0.74
(Wade, Bryant)	0.62

(b) The priority queue, including all intersection points between r and the x-axis which have been discovered between the y-axis and r. The grayed entry is the one being processed. The bolded entry is the one that is newly added by processing this point.

$\mathcal{P}$

Line	$m = 1$	$m = 2$
Durant	0.000	0.000
James	0.00	0.000
Wade	0.057	0.057
Nowitzki	0.062	0.059 (Wade)
Bryant	0.080	0.080
Anthony	0.089	0.084 (Bryant)
Stoudemire	0.105	0.066 (Wade)
Randolph	0.188	**0.083 (Wade)**

(c) The best cost values for each line as of the processing of the darker point in (a). The bold indicates the sustained value in the table after processing the point.

Figure 7: The processing of the tenth event point, the intersection of the lines corresponding to Stoudemire and Randolph. The distance along r of the intersection point from the contour is 0.034. In this case, the newly discovered length-2 chain, (Stoudemire, Randolph), has cost $max(0.105, 0.034)$, which does not improve on the value already found for Randolph, 0.083.

the best possible route along l_j to $p_{i,j}$. The cost of that route is the larger of the distance from $p_{i,j}$ to $\mathcal{C}_k$ and the cost of the best possible route to the previous intersection point of l_j. This value is updated for each of the m cells in row i.

On the other hand, for a chain that emerges along l_j, there are two possibilities, depending on the best route to get to $p_{i,j}$. Specifically, the best route of h turns to $p_{i,j}$ is the cheaper of the best route to the previous intersection point along l_j that used h turns and the best route to the previous intersection point along l_i that used $h - 1$ turns. The final cost is then the larger of the distance from $p_{i,j}$ to $\mathcal{C}_k$ and the minimum cost to $p_{i,j}$ as just described.

4.2 Algorithm Description

As we have hinted, the algorithm is a plane sweep through the arrangement of lines, searching for the minimal cost convex chain with fewer than c turns. The sweep features a ray r, originally positioned on the positive y-axis, moving radially through the positive quadrant to the positive x-axis. This plane sweep approach is appropriate as a consequence of Lemma 3.5, which reveals that the cost of any convex chain is maximized at an event point.

To initialize the algorithm, the intersection point of every line with the positive y-axis is processed, populating the data structures as in Figure 5. We only add to the priority queue the intersection points of lines that are immediate neighbours with respect to a sort on y-intercept, and for which the intersection point occurs in the positive quadrant. These points are maintained in the queue in descending order of angle from the positive x-axis (i.e., in the order in which the ray r will "sweep" them). The array $\mathcal{P}$ is initialised with empty paths for every cell, with a cost set to the distance of the relevant line to the contour.

The algorithm proceeds simply by popping the next event from $\mathcal{Q}$, updating the data structures as per Section 4.1.1, and pushing the new event points onto the queue. For the running basketball example, this is illustrated in Figure 7. For event points that correspond to vertices of the contour (since these, too, are intersection points of lines that will eventually be discovered by the plane sweep), we update every cell of $\mathcal{P}$ with new maximum costs for each line that has become more distant from $\mathcal{C}_k$.

Eventually, $\mathcal{Q}$ will be exhausted as r reaches the positive x-axis (see Figure 8). Every cell is updated with new maximum costs at this last contour vertex (the intersection of the contour with the x-axis). The final step is to scan through all of $\mathcal{P}$ and determine the smallest cost. This is the optimal solution, which is reported along with the path used to obtain it.

Algorithm 1 describes the algorithm with greater detail.

4.3 Asymptotic Complexity

THEOREM 2. *Algorithm 1 for the two-dimensional case finds an optimal k-regret minimizing set of order c in $\mathcal{O}(cn^2)$ time with $\mathcal{O}(n^2)$ space.*

PROOF. First consider space used. The size of the contour is bounded by n. Of the three data structures, $\mathcal{L}$ is of size exactly n, $\mathcal{P}$ is of size exactly $n * c \leq n^2$, and $\mathcal{Q}$ is proportional to the largest number of intersection points that have been discovered but not processed, clearly less than $n * (n - 1)$. Therefore, the total space is $\mathcal{O}(n^2)$.

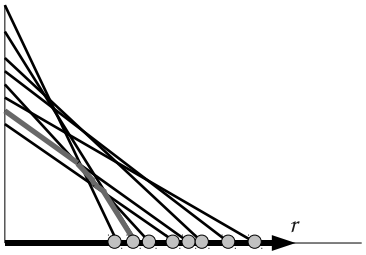
Regarding time, for each non-contour event point, of which there may be up to n^2, up to $2c$ cells of $\mathcal{P}$ are updated. For each contour event point, of which there are $|\mathcal{C}_k| \leq n$, each of the nc cells of $\mathcal{P}$ could potentially be updated. The initialisation requires a sort of n lines and then an initialisation of up to $n - 1$ insertions into $\mathcal{Q}$ and nc values for $\mathcal{P}$, which can be computed in constant time. At the conclusion of the plane sweep, all nc cells of $\mathcal{P}$ must be scanned. Therefore, the entire procedure is $\mathcal{O}(n^2c)$.

□

5. AN ALGORITHM FOR GENERAL DIMENSION

In Section 4, we gave an efficient plane sweep, dynamic programming algorithm to find optimal k-regret minimizing sets in two dimensions. Unfortunately, plane sweep algorithms do not readily generalize to higher dimensions. So, in this section, we offer a greedy algorithm which exploits Lemma 5.1 in order to progress towards an optimal solution.

LEMMA 5.1. *Consider the envelope produced by some set $\mathcal{L}$. For another set $\mathcal{L}'$ to better minimize regret, it is necessary that some line of $\mathcal{L}' \setminus \mathcal{L}$ either passes through the area between the contour and the envelope or remains entirely under the contour at the angle for which the distance ratio for S is maximized.*



(a) The arrangement of lines after processing every event point, emptying $\mathcal{Q}$.

	$\mathcal{P}$	
Line	$m = 1$	$m = 2$
Durant	0.114	0.114
James	0.208	0.208
Wade	0.625	0.625
Nowitzki	0.117	0.059 (Wade)
Bryant	0.566	0.566
Anthony	0.397	0.081 (Wade)
Stoudemire	**0.105**	**0.000 (Durant)**
Randolph	0.188	**0.000 (James)**

(b) The best cost values for each line at the conclusion of the plane sweep. The bold entries represent optimal solutions for each $j \leq m$.

Figure 8: The data structures at the conclusion of the plane sweep, when both r reaches the x-axis and $\mathcal{Q}$ is empty. Each cell (h, i) of $\mathcal{P}$ contains the optimum solution that contains i or fewer lines and ends with line l_h. The minimum value in the entire table is the optimal solution for this two-dimensional j-Regret problem.

PROOF. This results from Corollary 3.3, which indicates that for that fixed utility function, distance is directly proportional to the regret ratio. The L_2 distance metric is transitive. So, if no line in $\mathcal{L}'$ intersects the area underneath the lower envelope of $\mathcal{L}$ at the angle for which the distance ratio is maximized, then its maximum distance is clearly larger. □

We use the insight of Lemma 5.1 to design a greedy algorithm which behaves as follows. Consider a given set of lines S for which the distance ratio is maximized at point p. Consider also a new line $l \notin S$. We advance to a new set S' exactly when l passes between the origin and p and then if there is some element $l' \in S$ such that the maximum distance ratio of $S \setminus \{l'\} \cup \{l\}$ to $\mathcal{C}_k$ is less than that at p. That is to say, if l intersects the segment $[O, p]$, then we look for a new, better set that can be obtained by replacing some element of S with l. We know from Lemma 5.1 that the intersection test is a necessary condition to find some improved solution S'.

Overall, the algorithm begins with an initial seed solution and cycles through all lines repeatedly, conducting the above test in order to refine S, until no line can improve the cost any more. We report this solution. The greedy algorithm is feasible because, given the contour insight that we have developed in this paper, one can efficiently determine which of two sets is more optimal: it is the one for which the maximum distance of the lower envelope of the set from $\mathcal{C}_k$ is minimized. The algorithm is detailed in Algorithm 2.

Note that the algorithm is guaranteed to terminate because it will always progress towards a better solution due to Line 13, and there are finitely many subsets of $\mathcal{L}$.

6. RELATED WORK

The idea of representing an entire dataset with a few representative tuples for multi-criteria decision making has drawn much research attention in the past decade, since the introduce of the Skyline operator by Börzsönyi et al. [2]. However, its susceptibility to the curse of dimensionality is well-known. Chan et al. [4] made a compelling case for this, demonstrating that on the NBA basketball dataset (as it was at the time), more than 1 in 20 tuples appear in the skyline in high dimensions. Consequently, there have been numerous efforts to derive a smaller cardinality representative subset (e.g., [3, 13, 20]), especially one that presents very distinct tuples (e.g., [9, 16]).

Regret and regret minimizing sets are relatively new in the lineage of these efforts. When introduced by Nanongkai et al. [14], the emphasis was on proving that the maximum regret ratio is bounded by:

$$\frac{d-1}{(c-d+1)^{d-1}+d-1}.$$

Naturally, this bound holds for our generalisation introduced in this paper. As far as we know, no research has yet concerned itself with finding optimal regret minimizing sets.

The top-k query off which *regret* is based is well studied. The pareto-dominance graph [21] uses ideas of pareto-optimality to index for top-k queries. The Onion Technique [5] is a depth-based approach, but suffers the same curse of dimensionality as the skyline. Ilyas offers a nice survey on the topic of top-k queries [12]. Duality transforms are pervasive in this research area (e.g., [8, 17]). We use the duality transform of Chester et al. [6, 7] because of the immediate results on top-k depth contours it provides in answering the reverse top-k queries of Vlachou et al. [18, 19].

Using duality transforms on data points casts the problem into the context of arrangements. In two dimensions, plane sweep algorithms [10] are a typical approach to solving problems on arrangements of lines. Agarwal et al. [1] give bounds on the the number of edges and vertices that can exist at each level (or depth) of an arrangement. Top-k depth contours are not the only notion of contours or depth in arrangements of lines: Hugg et al. [11] evaluate several depth measures and Zuo et al. [22] derive general stastical results that apply to many of them and could be useful in extending this work. Rousseuw and Hubert consider depth in arrangements for dimensions greater than two [15], which could present deeper insight into the greey algorithm presented here.

7. CONCLUSION

Regret minimizing sets are a nice alternative to skyline as a succinct representative subset of a dataset, but suffer from a very strict assumption that users expect to see their top-1 choice for their queries. We generalised the notion to that of k-*regret minimizing sets*, which evaluates how representative a subset of a dataset is not by how closely it approximates every users' top-1 choice, but their top-k choice. We showed that in dual space, the top-k depth contour of a dataset is exactly the optimal k-regret minimizing set. If the cardinality of the k-regret minimizing set is specified as an input parameter, then the convex chain that minimizes the ratio of distances from itself to the contour and the contour to the origin is precisely the optimal solution. We used these ideas to construct an $\mathcal{O}(n^2c)$ optimal algorithm for two dimensions and a greedy algorithm for general dimension.

8. REFERENCES

[1] P. K. Agarwal, B. Aronov, and M. Sharir. On levels in arrangements of lines, segments, planes, and triangles. In *Proc. Symposium on Computational Geometry*, pages 30–38, New York, NY, USA, 1997. ACM.

[2] S. Börzsönyi, D. Kossmann, and K. Stocker. The skyline operator. In *Proc. International Conference on Data Engineering (ICDE)*, pages 421–430, Washington, DC, USA, 2001. IEEE Computer Society.

[3] C.-Y. Chan, H. V. Jagadish, K.-L. Tan, A. K. H. Tung, and Z. Zhang. Finding k-dominant skylines in high dimensional space. In *Proc. ACM Special Interest Group on Management of Data (SIGMOD)*, pages 503–514, New York, NY, USA, 2006. ACM.

[4] C.-Y. Chan, H. V. Jagadish, K.-L. Tan, A. K. H. Tung, and Z. Zhang. On high dimensional skylines. In *Proc. International Conference on Extending Database Technology (EDBT)*, pages 478–495, Berlin, Heidelberg, 2006. Springer-Verlag.

[5] Y.-C. Chang, L. Bergman, V. Castelli, C.-S. Li, M.-L. Lo, and J. R. Smith. The onion technique: indexing for linear optimization queries. In *Proc. ACM Special Interest Group on Management of Data (SIGMOD)*, pages 391–402, New York, NY, USA, 2000. ACM.

[6] S. Chester, A. Thomo, S. Venkatesh, and S. Whitesides. Indexing for vector projections. In *Proc. Database Systems for Advanced Applications (DASFAA)*, pages 367–376. Springer-Verlag, April 2011.

[7] S. Chester, A. Thomo, S. Venkatesh, and S. Whitesides. Indexing reverse top-k queries. 2012. arXiv:1205.0837v1 [cs.DB].

[8] G. Das, D. Gunopulos, N. Koudas, and N. Sarkas. Ad-hoc top-k query answering for data streams. In *Proc. Very Large Databases (PVLDB)*, pages 183–194. VLDB Endowment, 2007.

[9] A. Das Sarma, A. Lall, D. Nanongkai, R. J. Lipton, and J. Xu. Representative skylines using threshold-based preference distributions. In *Proceedings of the 2011 IEEE 27th International Conference on Data Engineering*, ICDE '11, pages 387–398, Washington, DC, USA, 2011. IEEE Computer Society.

[10] H. Edelsbrunner and L. J. Guibas. Topologically sweeping an arrangement. In *Proc.* 18^{th} *ACM Symposium on Theory of Computing*, pages 389–403, New York, NY, USA, 1986. ACM.

[11] J. Hugg, E. Rafalin, K. Seyboth, and D. Souvaine. An experimental study of old and new depth measures. In *In Proc. Workshop on Algorithm Engineering and Experiments (ALENEX06), Lecture Notes in Computer Science*, pages 51–64. Springer, 2006.

[12] I. F. Ilyas, G. Beskales, and M. A. Soliman. A survey of top-k query processing techniques in relational database systems. *ACM Compututing Surveys*, 40(4):11:1–11:58, Oct. 2008.

[13] J. Lee, G.-w. You, and S.-w. Hwang. Personalized top-k skyline queries in high-dimensional space. *Information Systems*, 34(1):45–61, Mar. 2009.

[14] D. Nanongkai, A. D. Sarma, A. Lall, R. J. Lipton, and J. Xu. Regret-minimizing representative databases. volume 3, pages 1114–1124. VLDB Endowment, 2010.

[15] P. J. Rousseeuw and M. Hubert. Depth in an arrangement of hyperplanes, 1999.

[16] Y. Tao, L. Ding, X. Lin, and J. Pei. Distance-based representative skyline. In *Proc. International Conference on Data Engineering (ICDE)*, pages 892–903, Washington, DC, USA, 2009. IEEE Computer Society.

[17] P. Tsaparas, N. Koudas, and T. Palpanas. Ranked join indices. In *Proc. International Conference on Data Engineering (ICDE)*, pages 277–288, 2003.

[18] A. Vlachou, C. Doulkeridis, Y. Kotidis, and K. Norvag. Reverse top-k queries. In *Proc. International Conference on Data Engineering (ICDE)*, pages 365–376. IEEE, March 2010.

[19] A. Vlachou, C. Doulkeridis, Y. Kotidis, and K. Nørvåg. Monochromatic and bichromatic reverse top-k queries. *Transactions on Knowledge and Data Engineering (TKDE)*, 23(8):1215–1229, 2011.

[20] M. L. Yiu and N. Mamoulis. Multi-dimensional top-k dominating queries. 18(3):695–718, June 2009.

[21] L. Zou and L. Chen. Pareto-based dominant graph: An efficient indexing structure to answer top-k queries. *IEEE Trans. Knowl. Data Eng.*, 23(5):727–741, 2011.

[22] Y. Zuo and R. Serfling. Structural properties and convergence results for contours of sample statistical depth functions. *Annals of Statistics*, 28:483–499, 2000.

APPENDIX

A. ALGORITHM PSEUDOCODE

Algorithm 1 Calculating an optimal k-max-regret minimizing set S from $D \subseteq \mathbb{R}^2$ with $|S| \leq c$

Input: $\mathcal{C}_k$; c; $\mathcal{L}$, sorted by y-intercept
Output: $S \subseteq \mathcal{L}$, the lines that together form an optimal solution S with $|S| \leq c$
if $|\mathcal{C}_k| \leq c$ **then**
 Return $\mathcal{C}_k$
end if
Initialize $\mathcal{Q}$ as an empty priority queue; priority is angle of points, desc
for all $l \in \mathcal{L}$ **do**
 Set $\mathcal{P}(l) = k * [(y\text{−intercept}, \max(y\text{−intercept} - y\text{−interceptof } \mathcal{C}_k, 0)]$
 Add to $\mathcal{Q}$ intersect(l and next l) if not last l
end for
while $\mathcal{Q}$ is not empty **do**
 Let p be next point in $\mathcal{Q}$
 Let Δ be distance ratio of p to $\mathcal{C}_k$ and $\mathcal{C}_k$ to O
 if $p \in \mathcal{C}_k$ **then**
 for all ($v \in \mathcal{P}$) **do**
 Let Δ' be distance of ratio of line to $\mathcal{C}_k$ and $\mathcal{C}_k$ to O
 Let v be $max(v, \Delta')$
 end for
 end if
 Retrieve adjacent l_i, l_{i+1} that intersect at p
 Add intersect(l_{i-1}, l_{i+1}) if angle less than that of p
 Add intersect(l_i, l_{i+2}) if angle less than that of p
 for all $j \in [0, c)$ **do**
 Let $\mathcal{P}(l_i)_j = \max(\mathcal{P}(l_i)_j, \Delta)$
 if $j \geq 0$ **then**
 if $\mathcal{P}(l_{i+1})_{j-1} < \mathcal{P}(l_{i+1})_j$ **then**
 Add p to path of $\mathcal{P}(l_{i+1})_j$
 Let $\mathcal{P}(l_{i+1})_j = \max(\mathcal{P}(l_{i+1})_{j-1}, \Delta)$
 else
 Let $\mathcal{P}(l_{i+1})_j = \max(\mathcal{P}(l_{i+1})_j, \Delta)$
 end if
 else
 Let $\mathcal{P}(l_{i+1})_j = \max(\mathcal{P}(l_{i+1})_j, \Delta)$
 end if
 end for
 Swap l_i and l_{i+1}
end while
for all $l \in \mathcal{L}, j \in [0, c)$ **do**
 Remember $v = \mathcal{P}(l)_j$ if smallest yet seen, breaking ties with smaller j
end for
RETURN set of lines generating path corresponding to v

Algorithm 2 Greedy algorithm to compute k-max-regret minimizing set S from $D \subseteq \mathbb{R}^d$ with $|S| \leq m$

Input: $\mathcal{C}_k$; m; $\mathcal{L}$
Output: $S \subseteq \mathcal{L}$, the lines that together form a solution S with $|S| \leq m$
if $|\mathcal{C}_k| \leq m$ **then**
 Return $\mathcal{C}_k$
end if
Select an arbitrary set $S \subseteq \mathcal{L}$, such that $|S| = m$
Let p be point at which distance ratio from lower envelope of S to $\mathcal{C}_k$ is maximized
Place all lines $l_i \notin S$ into unsorted queue, $\mathcal{Q}$.
while $\mathcal{Q}$ is not empty **do**
 Let l be next line in $\mathcal{Q}$
 if l intersects $[O, p]$ **then**
 for all ($l' \in S$) **do**
 if Distance ratio of $S \setminus \{l'\} \cup \{l\}$ to $\mathcal{C}_k <$ distance ratio of S to $\mathcal{C}_k$ **then**
 Let S be $S \setminus \{l'\} \cup \{l\}$
 Restore to $\mathcal{Q}$ all $l_i \notin S$
 Break
 end if
 end for
 end if
end while
RETURN S